

REMARKS

Claims 1-3, 5-7, and 9-22 are pending in the present application. Claims 4 and 8 are canceled. Claims 1, 2, 5, 7, 12, 13, 16, and 18-20 are amended. The amendments are supported by the specification at least on page 8, line 7, to page 11, line 22, as seen in the description of the illustrative embodiments, as well as the examples presented in the disclosure. Claims 21 and 22 are added. Reconsideration of the claims is respectfully requested.

Amendments are made to the specification to correct errors and to clarify the specification. No new matter is added by any of the amendments to the specification.

I. Examiner Interview

A telephone interview was conducted with the Examiner on May 25, 2005. Applicants thank the Examiner for the courtesies extended during the interview. Applicants' representative proposed several amendments to the claims. The Examiner agreed to reconsider the rejection, but no agreement was reached with regard to allowable subject matter.

II. 35 U.S.C. § 102, Anticipation

The Office Action rejects claims 1-20 under 35 U.S.C. § 102 as being anticipated by *Padmanabhan et al.* (U.S. Pub. No. 2004/0107300 A1). This rejection is respectfully traversed.

Padmanabhan teaches defining a default pathset, such as "all available paths," for a host to a device. See *Padmanabhan*, paragraph [0134]. *Padmanabhan* also teaches that a driver checks the pathset information assigned to a packet and, if none has been defined at the packet level, the driver uses the default pathset that is defined for the device. See *Padmanabhan*, paragraph [0135]. In other words, only one set of paths, the default pathset, is defined for a device. Also, *Padmanabhan* teaches that if a path fails, the host may failover to another path. See *Padmanabhan*, paragraph [0149]. However, *Padmanabhan* does not teach how the failover path is selected. Presumably, the failover path is another path in the default pathset.

In contradistinction, the present invention performs multiple path input/output where a set of primary paths to be used for issuing transactions to the device and a set of standby paths to be used for failover are selected. The set of primary paths is a first subset of all available paths for the device and the set of standby paths is a second subset of all available paths for the device, where the first subset and the second subset have no paths in common. A path control module for the device is configured with the set of primary paths and the set of standby paths. Transactions are issued to the device using the set of primary paths. When a path in the set of primary paths fails, the present invention fails over using at least one path in the set of standby paths. *Padmanabhan* does not teach or fairly suggest configuring a path control module with two disjoint sets of paths, one for issuing transactions and one for failover. More particularly, *Padmanabhan* does not teach configuring a path control module with a set of primary paths and configuring the path control module with a set of standby paths where each set of paths is a subset of all available paths.

The applied reference does not teach each and every claim limitation; therefore, *Padmanabhan* does not anticipate claim 1. Independent claims 12 and 18 recite subject matter addressed above with respect to claim 1 and are allowable for similar reasons. Since claims 2, 3, 5-7, 9-11, 13-17, 19, and 20 depend from claims 1, 12, and 18, the same distinctions between *Padmanabhan* and the invention recited in claims 1, 12, and 18 apply for these claims. Additionally, claims 2, 3, 5-7, 9-11, 13-17, 19, and 20 recite other additional combinations of features not suggested by the reference.

More particularly, with respect to claim 2, the Office Action alleges that *Padmanabhan* teaches failing over to the set of standby paths at paragraph [0018], which reads as follows:

[0118] In the example discussed above for prior systems, where the driver retries some number of times (e.g. twice) before sending a failure message up to the metadriver layer, in the present invention the vHCI can immediately (after a single failure) fail over to another path. Thus, the new system requires only two tries (one failed and one successful) to complete the I/O request, rather than four tries for the example given for prior systems, resulting in a significant time savings.

This portion merely states that the vHCI can fail over to another path; however, *Padmanabhan* provides no teaching with regard to how to select the failover path. Certainly, *Padmanabhan* does not teach a set of standby paths or configuring a path control manager to issue transactions to a set of standby paths responsive to a path in the set of primary paths failing, as recited in claim 2. Therefore, *Padmanabhan* does not anticipate claim 2.

Furthermore, claim 2 is amended to include a limitation originally presented in claim 4. With respect to claim 4, the Office Action alleges that *Padmanabhan* inherently teaches marking a first path as down or inactive because *Padmanabhan* allegedly shows information about an available path at paragraph [0160], which reads as follows:

[0160] When discovery is complete, an I/O request coming down from a target driver via the common architecture layer 270 is sent to the vHCI layer 280. The vHCI provides the requested device information to the MDI, and the MDI—which has the information about paths to the devices—selects and sends back information about an available path to the vHCI.

The Office Action misapplies the concept of “inherent” anticipation. Section 102 of Title 35 deals with novelty and loss of patent rights. An invention is said to be “anticipated” when it is squarely described or disclosed in a single reference as identified from one of the categories of 35 U.S.C. § 102, commonly referred to as “prior art.” Express anticipation occurs when the invention is expressly disclosed in the prior art, patent or publication.

In some cases, however, when the claimed invention is not described *in haec verba*, the “doctrine of inherency” is relied on to establish anticipation. Under the principles of inherency, a claim is anticipated if a structure in the prior art necessarily functions in accordance with the limitations of a process or method claim. *In re King*, 801 F.2d 1324, 231 U.S.P.Q. 136 (Fed. Cir. 1986). A prior art reference that discloses all of a patent’s claim limitations anticipates that claim even though the reference does not expressly disclose the “inventive concept” or desirable property the patentee discovered. *Verdgaal Brothers, Inc. v. Union Oil Company of California*, 814 F.2d 628, 2 U.S.P.Q.2d 1051, (Fed. Cir. 1987). It suffices that the prior art process inherently possessed at that property. *Id.* Mere possibilities or even probabilities, however, are not enough to

establish inherency. The missing claimed characteristics must be a "natural result" flowing from what is disclosed. *Continental Can Co. v. Monsanto Co.*, 948 F.2d 1264, 20 U.S.P.Q.2d 1746 (Fed. Cir. 1991). Unstated elements in a reference are inherent when they exist as a "matter of scientific fact." *Constant v. Advanced Micro-Devices, Inc.*, 848 F.2d 1560, 7 U.S.P.Q.2d 1057 (Fed. Cir.), *cert. denied*, 488 U.S. 892 (1988) and *Hughes Aircraft Co. v. United States*, 8 U.S.P.Q.2d 1580 (Ct. Cl. 1988). Otherwise, the invention is not inherently anticipated.

In the present case, the assertion that these elements are present can be made only through the use of Applicants' disclosure as a template to fill in the missing elements. The mere teaching of a discovery operation would not lead a person of ordinary skill in the art to conclude that the vHCI of *Padmanabhan* necessarily marks a path as down or inactive. Presumably, the vHCI of *Padmanabhan* would retain the path in the default pathset and retry the path at another time. In other words, the claimed features are not a natural result or matter of scientific fact in the applied reference. Rather, the disclosure of *Padmanabhan* is clearly deficient and the Office Action fails to establish a *prima facie* case of anticipation for claim 2, as amended. Claims 13 and 19 recite subject matter addressed above with respect to claim 2 and are allowable for similar reasons.

With respect to claim 3, the Office Action alleges that *Padmanabhan* teaches determining a second path within the set of standby paths and issuing a transaction, for which a primary path has failed, to the second path at paragraph [0118], which is shown above, and paragraphs [0121] and [0122], which read as follows:

[0121] A failover operation in the embodiment of FIGS. 3-4 proceeds as follows. When a given path such as path 292 (to bus 296, which connects to devices 330) fails, another path (e.g. path 294) is needed. It would be possible to code all the different device-specific information at the vHCI layer 280, but this would mean that any new device type that is added would require modification of the vHCI layer. Thus, preferably a set of one or more failover ops (operations) modules 272-276 is created, once for each type of storage device.

[0122] When the vHCI 280 needs to activate a path, it accesses the appropriate failover ops module (e.g. module 272) and sends an "activate" command. The module 272

then connects to the appropriate HBA (pHCI) driver with the correct protocol.

The cited portion makes no mention whatsoever of a set of secondary paths.

Padmanabhan simply fails to teach determining a path within a set of standby paths and issuing a request to the second path. Claim 14 recites subject matter addressed above with respect to claim 3 and is allowable for similar reasons.

With respect to claim 5, the Office Action alleges that *Padmanabhan* teaches responsive to a first failed path being restored, failing back to the set of primary paths at paragraph [0146], which reads as follows:

[0146] Enable multipath device configurations to dynamically self-assemble during boot and dynamic reconfiguration, not relying upon on-disk configuration databases to describe the multipath configuration.

Neither the cited portion nor any other portion of *Padmanabhan* teaches or suggests failing back to the set of primary paths responsive to a first path being restored.

Padmanabhan does not contemplate issuing transactions to any other set of paths for failover. *Padmanabhan* also makes no mention whatsoever of a path being restored. The Office Action appears to cite an arbitrary portion of the reference that has nothing to do with a restored path or failing back to a set of primary paths. Thus, the Office Action fails to establish that claim 5 is *prima facie* anticipated. Claim 15 recites subject matter addressed above with respect to claim 5 and is allowable for similar reasons.

With respect to claim 7, the Office Action alleges that *Padmanabhan* teaches determining a second path within the set of standby paths responsive to a transaction failing at paragraph [0018], which is shown above. This cited portion only teaches failing over to "another path." The Office Action proffers no analysis as to why this is somehow equivalent to selecting a path within a set of standby paths and issuing a transaction to the second path. In fact, *Padmanabhan* does not even teach a set of standby paths. As such, the Office Action fails to establish a *prima facie* case of anticipation for claim 7. Claims 16 and 20 recite subject matter addressed above with respect to claim 7 and are allowable for similar reasons. Claims 7, 16, and 20 also recite marking the failed path as down or inactive. *Padmanabhan* fails to teach or suggest this feature, as discussed above with respect to claim 2.

With respect to claim 9, the Office Action alleges that *Padmanabhan* teaches adding the first path back to the set of primary paths and adding the second path back to the standby paths when the first path is restored, at paragraph [0146], which is shown above. Again, *Padmanabhan* does not contemplate issuing transactions to any other set of paths for failover. *Padmanabhan* also makes no mention whatsoever of a path being restored. The Office Action appears to cite an arbitrary portion of the reference that has nothing to do with a restored path or returning paths to a set of primary paths or a set of standby paths. In fact, *Padmanabhan* makes no mention of a set of standby paths. Thus, the Office Action fails to establish that claim 9 is *prima facie* anticipated. Claim 17 recites subject matter addressed above with respect to claim 9 and is allowable for similar reasons.

Still further, with respect to claim 11, the Office Action alleges that *Padmanabhan* teaches configuring a first path control module for a first device with a first set of primary paths and a first set of standby paths and configuring a second path control module for a second device with a second set of primary paths and a second set of standby paths, wherein the set of standby paths for the second device is the set of primary paths for the first device and wherein the set of standby paths for the first device is the set of primary paths for the second device, at paragraph [0149], which reads as follows:

[0149] Support automatic failover to route I/O requests through alternate active paths on transport failures.

Clearly this cited portion makes no mention of two path control modules for two devices with two sets of primary paths and two sets of secondary paths. This cited portion also makes no mention of a set of primary paths for one device being the same set of paths as the set of standby paths for another device. The Office Action appears to cite an arbitrary portion of the reference and simply conclude that the features are taught. The Office Action proffers no analysis whatsoever as to why failover from one path to another is somehow equivalent to two path control modules for two devices with two sets of primary paths and two sets of standby paths. The Office Action fails to establish a *prima facie* case of anticipation for claim 11.

New claims 21 and 22 recite that the path control module is a dynamically loaded extension of a device driver for the device. *Padmanabhan* does not teach or suggest this feature.

Therefore, Applicants respectfully request withdrawal of the rejection of claims 1-3, 5-7, and 9-20 under 35 U.S.C. § 102.

Furthermore, *Padmanabhan* does not teach, suggest, or give any incentive to make the needed changes to reach the presently claimed invention. *Padmanabhan* actually teaches away from the presently claimed invention because it teaches a single pathset for a device, as opposed to a set of primary paths for issuing transactions and a set of standby paths for failover, as in the presently claimed invention. Absent the Office Action pointing out some teaching or incentive to implement *Padmanabhan* to configure a path control module with a set of standby paths for failover, one of ordinary skill in the art would not be led to modify *Padmanabhan* to reach the present invention when the reference is examined as a whole. Absent some teaching, suggestion, or incentive to modify *Padmanabhan* in this manner, the presently claimed invention can be reached only through an improper use of hindsight using Applicants' disclosure as a template to make the necessary changes to reach the claimed invention.

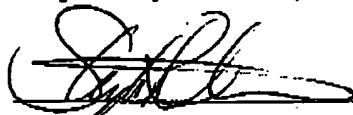
III. Conclusion

It is respectfully urged that the subject application is patentable over the prior art of record and is now in condition for allowance.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: June 22, 2005

Respectfully submitted,



Stephen R. Tkacs
Reg. No. 46,430
Yee & Associates, P.C.
P.O. Box 802333
Dallas, TX 75380
(972) 385-8777
Agent for Applicants